The XY Map Display Program

Matthew Marcus

August, 2002

I. Introduction

This program is designed to allow the user to look at and analyze fluorescence maps produced by the XY mapping program on 10.3.2. Such maps are multichannel images of a sample with the different channels representing the count rates seen in different energy bands. The program has the following facilities:

1. Input and normalization

Count rates can be normalized to the time per point, the incident intensity as measured at the end of each line, or not normalized at all. Additionally, the incident intensity sometimes decreases smoothly, in which case it can be fit to an exponential decay in order to reduce noise.

2. Single-channel display

The main screen's most obvious feature is an image area in which one can display a grayscale image of any channel. The white level and gamma of this image are under user control. The image may be zoomed and values read off with a crosshair. The image may also be saved as a simple spreadsheet array.

3. Tricolor display

You can pick three channels to be represented by red, green and blue. The white and black levels of each are adjustable. The resulting image can be saved as a bitmap or printed. If you make all three colors correspond to the same channel, you get a grayscale image. You can flip a switch and have it come out as a negative, which is often a better way to display the data than a positive image.

4. Cross-correlation

This is a subprogram which lets you look at the cross-correlation function between any two channels. It may be used to quantify how tightly associated one element is with another, and it also shows some information about the size scale of features in the sample.

5. Scatterplot

This function lets you plot the counts in one channel against those in another. It's part of the same subprogram as the cross-correlation plot. A scatterplot can help you identify distinct populations in a sample, for instance places where Zn is associated with Mn vs. places where Mn occurs alone.

II. Entry point – main screen

When you start the program, it asks for a file, which would be a `.xrf` file as produced by the XY mapping program. It then shows you the main screen, as shown in Figure 1. From here, you can adjust the display or enter the various sub-programs.

When you start, the image is often not visible. This is because the XY data range on the image display is not set for the current data. Use the graph tools (X,Y rescale buttons) at the bottom of the image to reset. The program automatically sets the aspect ratio so that the display is isotropic, that is 1mm on X takes up the same amount of screen space as 1mm on Y. Note that X goes to lower numbers at the right, and Y decreases going down. These directions match the way the data are taken.

The image display is a positive grayscale, so black corresponds to small amounts of the selected element and white to large amounts. To select the channel of interest, use the ROI selector at upper left. This is a digital control linked to the `SCA Description` indicator just under the `File path` indicator. This indicator shows the channel range and the name for the selected ROI, as determined in the XY mapping program.

The display controls include the gamma slider at the top, the `Auto/Man`ual scaling button at the right, the color box just above the `Auto/Man` button, and the white-level slider running down the right side. When gamma is all the way to the left, the display emphasizes the lowlights relative to the brighter areas. Similarly, when gamma is to the right, the bright spots are made brighter relative to the background.

The color box lets you tint the grayscale image to suit your tastes. It makes the high-intensity pixels come out in whatever color you select, while the low-intensity pixels are black.

When auto scaling is enabled by the `Auto/Man` button, the white level is set to the highest intensity in the image. This often results in an image consisting of a few

white dots on a black ground, due to the presence of 'hotspots' in the sample. To make the rest of the image visible, push the button so that it shows `Man` instead of `Auto`, and use the scale slider. You may have to change the top level on this slider; this is done by double-clicking on the top number and entering a new one.

As stated above, there are four normalization modes:

`None`: Just reports the counts, as in the XY mapping program.

`Dwell time`:
Divides the counts by the dwell time, yielding the count rate.

`Dwell time and I0`:
Divides the counts by the product of dwell time and I_0 reading, thus yielding a normalized intensity. This is the default mode.

`Dwell time and fitted I0`:
Does an exponential fit to I_0 as a function of time, then divides the data by the product of this fit and the dwell time. Useful only when such a fit works, which in the present example is not so.

The history of I_0 vs. line number (hence, time) is shown in the I_0 plot to the left. The white points are data, taken at the end of each line, and the red line is the exponential fit.

The graph and cursor tools are the ones provided with any LabVIEW graph. To get a reading of the intensity at any point, make sure the cursor is locked to the plot (the small lock icon at the right end of the cursor toolbar). Click on this if the lock is open. The zoom tools can be used to look at any area. The tricolor program (see below) will then work just on the visible area.

The `Save` button lets you save the image as a simple, 2D spreadsheet array, thus stripping off all header information. This might be useful for interfacing with other programs such as Excel.

The round green buttons below the center of the image display let you into the other screens.

III. Tricolor

When you push the Tricolor button, you get a screen as is shown in Figure 2. The top part of the Figure shows an overall view of the screen, while the bottom part shows just the controls. The basic idea is that the intensities in the Red, Green and Blue

channels are to be controlled by the counts in each of three channels. Each of these channels may be auto-scaled or have its black and white levels set manually. Thus, if all three channels are mapped to the same element and with the same scaling, you get a grayscale image as in the main screen. Unlike the main screen, this image can be saved, printed or inverted (negative contrast).

The main part of the set of controls is devoted to the selection and scaling of the channels. On the left are listed the names of all the channels. Next come the black and white levels, which are automatically set on entry. Next is a row of three white buttons. We thus have a 3xN matrix of buttons, where N is the number of channels. This matrix represents the selection of which color goes with which channel. Pushing a button in the first column sets the Red intensity to represent a particular channel, while buttons in the middle column set Green and those in the third column set the Blue. In the current example, we have Red mapped to Fe, Green to Mn and Blue to Zn. Thus, in the image, we can see that Zn is present mostly in the outermost layers, which appear blue. Fe and Mn alternate in the interior of this sample, as shown in by the relatively pure red and green colors.

You can suppress any of the colors by clicking its button again. Thus, you can make a bicolor Fe/Mn map by clicking the Blue button again to un-set it. The rule is that up to one button in each column may be set.

Just as you usually need to rescale the intensities in the single-element plot in the main screen, you usually will want to adjust the levels in the tricolor plot. To do this, flip the little `Autoscale?` switch at the right side of the control cluster, then enter a new value for the white or black level. Alternatively, you can use the green buttons just above the `RGB Select` button matrix to raise or lower the sensitivity (adjust white level) for any of the colors. As with manual entry, you have to flip the `Autoscale?` switch to make this work. These buttons multiply or divide the current white level by a factor set by the `Scale Jump` control at the bottom.

When you do a grayscale image, it's often useful to have it come out in negative contrast, that is with the higher intensity showing as dark. Such a negative image works better on the printed page and often shows subtle detail better than a positive image. Use

the Image Sign switch to make this happen. Try it on a tricolor image for some interesting visual effects.

You can save or print the image once you get it the way you like. The image may be saved in `.bmp` or `.jpg` format, selected by the control next to the `Write File` button. You can also print the image. To select the printer, select `File→Page Setup→Printer setup` from the menu bar at the top of the screen. This is a LabVIEW function and so may be done from any LabVIEW program.

When you're done with tricolor imaging, hit the red `Return` button to get back to the main screen. A common error is to minimize the tricolor screen, then try to do something with the main screen which is still visible behind the tricolor program. This won't work because control is still in the tricolor screen until you hit the `Return` button. The same problem occurs with the other sub-programs.

IV. Correlation

The cross-correlation between two signals I_1 and I_2 is defined as:

$$\rho(\vec{x}) \equiv \frac{< (I_1(\vec{y}-\vec{x}) - \overline{I}_1)(I_2(\vec{y}) - \overline{I}_2) >_{\vec{y}}}{\sqrt{< (I_1(\vec{y}) - \overline{I}_1)^2 >_{\vec{y}} < (I_2(\vec{y}) - \overline{I}_2)^2 >_{\vec{y}}}} \tag{1.1}$$

where $\overline{I}_1$ is the mean of I_1 over the sample area and $< \cdots >_{\vec{y}}$ is the mean of the quantity inside the brackets. Now, for $\vec{x} = 0$, the cross-correlation $\rho(\vec{x})$ reduces to the Pearson r between the two quantities, taken pixel-by-pixel. This correlation is a quantitative measure of the degree of association between two elements. The spatial scale of the correlation gives a measure of the size of features resolved in the sample. You can make a rough estimate of whether the correlation is statistically significant by comparing the amplitude of the feature at $\vec{x} = 0$ with those at larger distances, which are presumably due to noise or variations within the sample.

Often, the sample does not cover the whole scan area. This is true in the example shown. The area around the sample is essentially filled with zeros. This effect causes artifacts in the correlation plot which show up as horizontal and vertical streaks. To get rid of this artifact, we want to pad the non-sample area with a quantity equal to the mean of the signal within the sample. To do this, we must distinguish between sample and

non-sample. In this software, this distinction is done with a simple thresholding procedure. Only pixels with a signal intensity greater than a given amount are considered to be part of the sample. This procedure is done separately for each of the two signals involved.

The correlation plot is accessed by pressing the `Correlation` button shown in Figure 1. The resulting screen is shown in Figure 3. The signals to be correlated are selected using the green controls `SCA1` and `SCA2`. The `Threshold` controls are on the left side. The small black&white images are binary plots showing which pixels are included for each signal. The controls are normalized so that 1 is the maximum. Thus, no pixels will be included if the `Threshold` control is set to 1. Typical settings are 0.05-0.2.

The large blue intensity graph to the right shows the cross-correlation function. Zoom and crosshair tools are provided. The correlation plot may be saved in a format which looks like a `.xrf` file (actually, it gets a `.cor` extension) so may be read in again by the XY map program. The plot may be saved as a bitmap or printed, just as in the tricolor sub-program.

The scale bar to the right of the plot shows the color scale and the maximum and minimum values found. The value at the center of the plot (Pearson r) is usually one of the extrema so may be read off the scale bar.

V. Scatterplot

The correlation graph is on one tab of a tab control. The other tab has the scatterplot, an example of which is shown in Figure 4. This is the scatterplot corresponding to the correlation plot shown in Figure 3, and done with the same thresholding. The effect of the thresholds may be seen as the sharp cutoffs of the plot on the left and bottom.

As with the other plots, the graph may be printed and saved. A cursor may be activated with the `Cursor` switch at the bottom. Also, the graph axes may be set to autoscale or manual via the appropriate switch, next to the `Cursor` switch.

The scatterplot gives you an idea of how associated the two elements are, and whether there are multiple populations, each with its own pattern of association. For

instance, one might argue that the tricolor image shows that the outer layers of the sample have a different ratio of Zn to Mn than the inner ones, and therefore constitute a different population, perhaps a different set of minerals. In this case, the scatterplot shows a line with a fringe of points extending up (higher Zn/Mn) from it. One might wonder if these two features correspond to different regions in the sample. To answer this question, we would want to isolate the points corresponding to the fringe and look at them separately. This is what the masking function is for.

VI. Masking

Every pixel in the map corresponds to a point on the scatterplot. Thus, a set of points on the scatterplot maps into a set of points on the map image. By selecting a region on the scatterplot, we select some set of points on the map. This selection is what is referred to as masking. In the present example, we want to know if the fringe of points seen in the Zn/Mn scatterplot comes from any specific part of the sample. We can define the set of points by drawing a mask as shown in Figure 5. The procedure is:

1. Turn on the cursor with the `Cursor` switch.
2. Position the cursor where you want the first point to be and push the `Add Pt to Mask` button.
3. Place the cursor on the next point and push the `Add Pt to Mask` button again. A line will be drawn on the screen.
4. Repeat until you have drawn all points in the polygon.
5. Push the `Complete Mask` button. This will close the polygon but will not draw that last side.
6. If you make an error, hit the `Cancel Mask` button and start over.

Now a mask has been defined, but there's no obvious effect. To see the effect on the scatterplot or correlation, flip the `Mask Display Mode switch`. The threshold binary plots will then reflect the masking, as will the correlation and scatterplots. Next, hit the `Return` button to get back to the main screen. The mask lines will appear on the images produced by the `Save` and `Print` buttons.

Now, in the main screen, push the `Use Mask` button and those areas not covered by the mask will go black. This effect is shown in Figure 6. The masking will apply to

all subsequent functions. Push the `Ignore Mask` button to see the whole sample again. Push the `Clear Mask` button to remove the mask altogether. Note that pushing the Ignore Mask button does not remove the mask, but affects only the display in the main screen. If you have a mask in effect and make another one by going back into the scatterplot routine, the result will be the intersection of the two masks. This might be useful for isolating sub-populations.

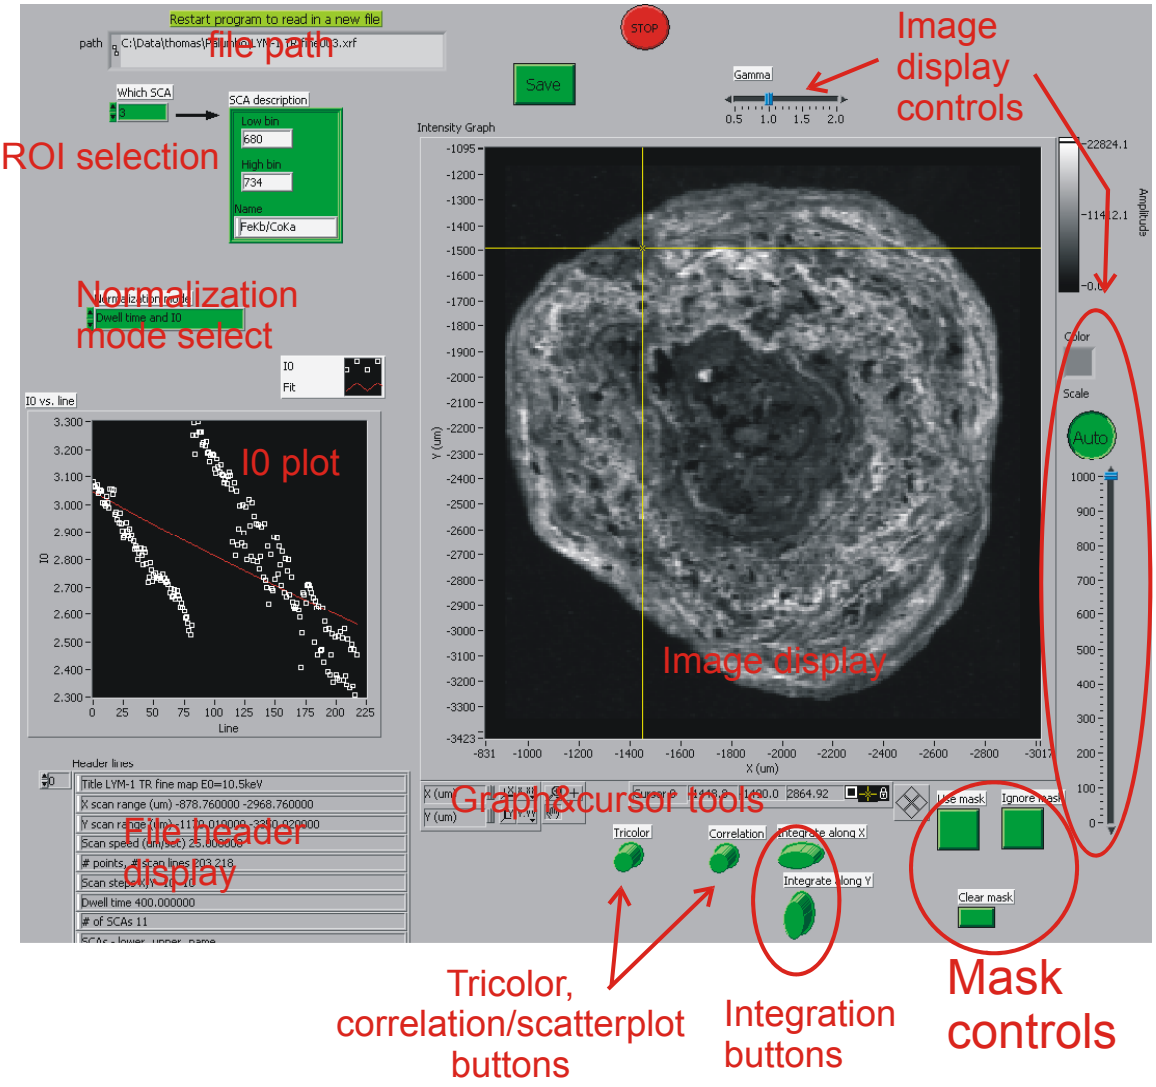

Figure 1. The main screen, with controls and groups of controls labeled in red. We are now looking at Fe, using the Kβ line to avoid interference from Mn. The normalization mode is the default, `Dwell time and I0`.

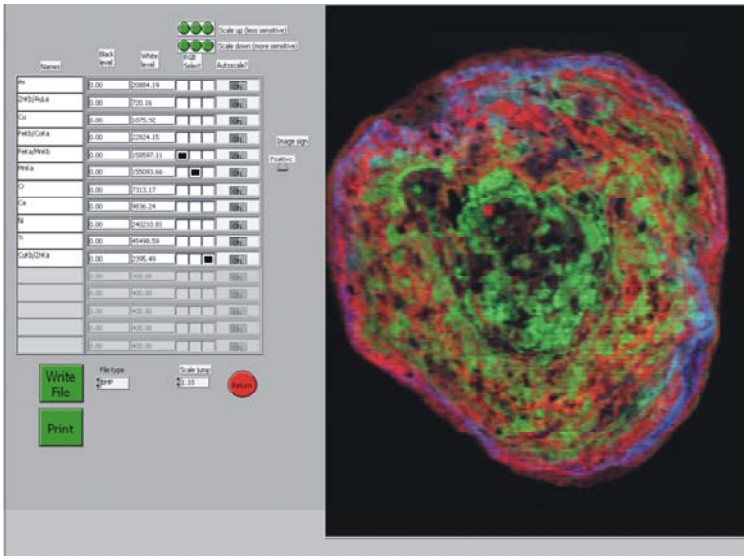

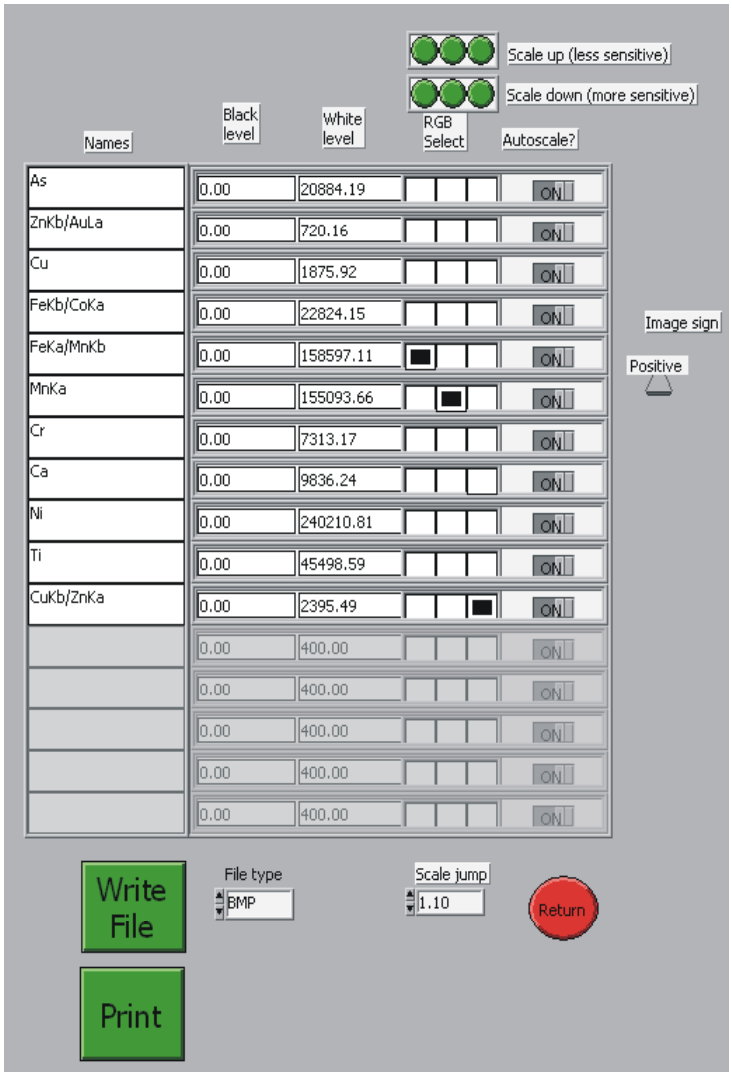

Figure 2. The tricolor screen. The whole screen is shown at the top, while an expanded view of the controls is shown below. In this example, Fe is mapped to red, Mn to green and Zn to blue. Autoscaling is used for all channels.

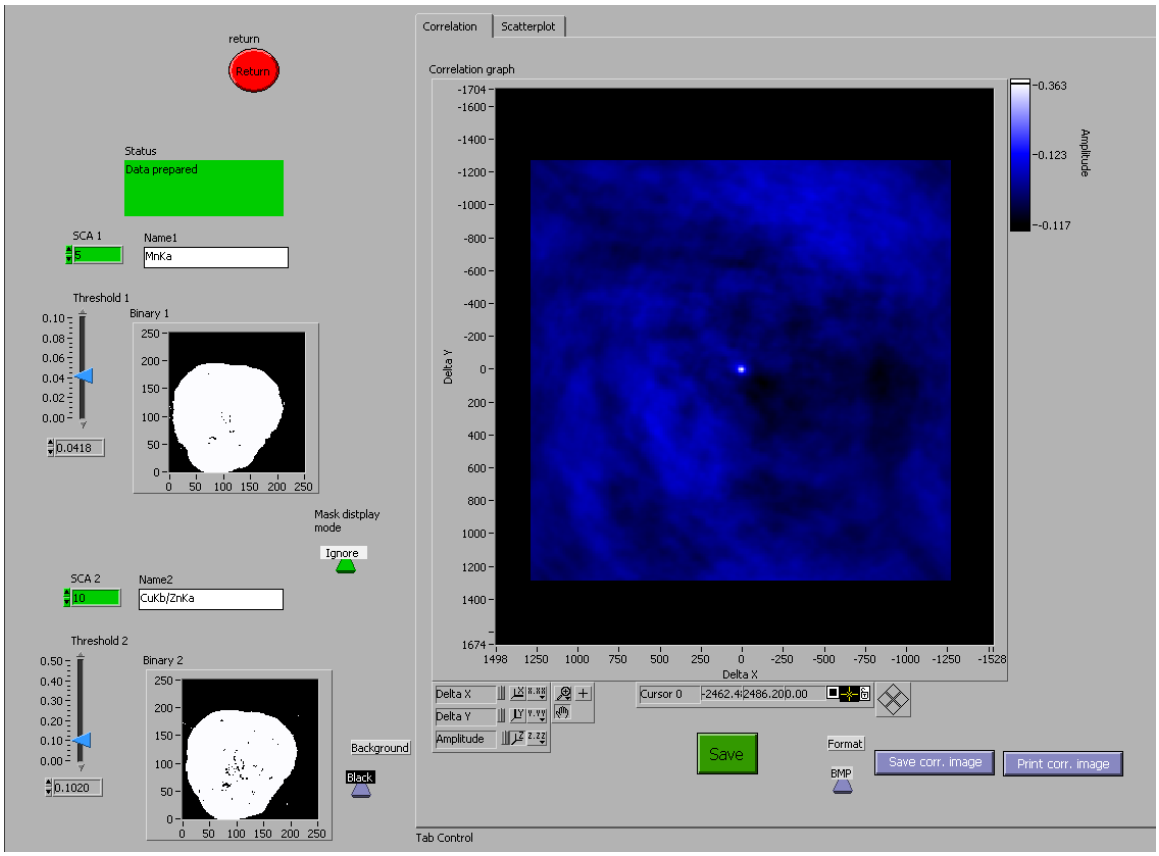

Figure 3. The `Correlation` screen. This is showing the correlation between Zn and Mn. The bright feature at the center demonstrates a significant association between these two species. The `Threshold` controls are set so that the only pixels within the sample are considered.

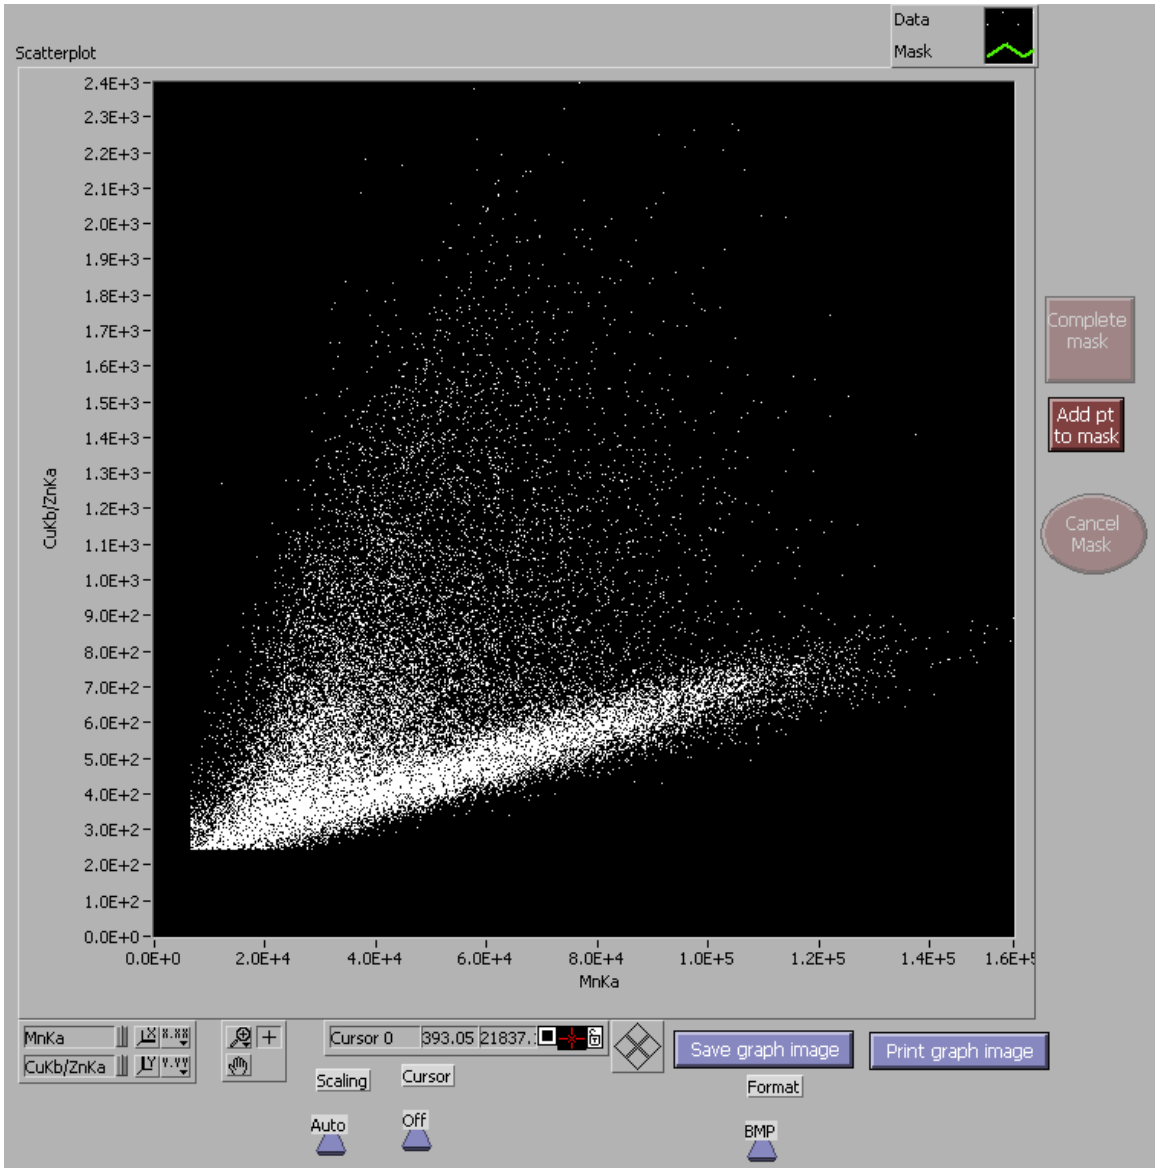

Figure 4. The scatterplot corresponding to the correlation shown in Figure 3. The sharp cutoffs at the left and bottom are due to the thresholding.

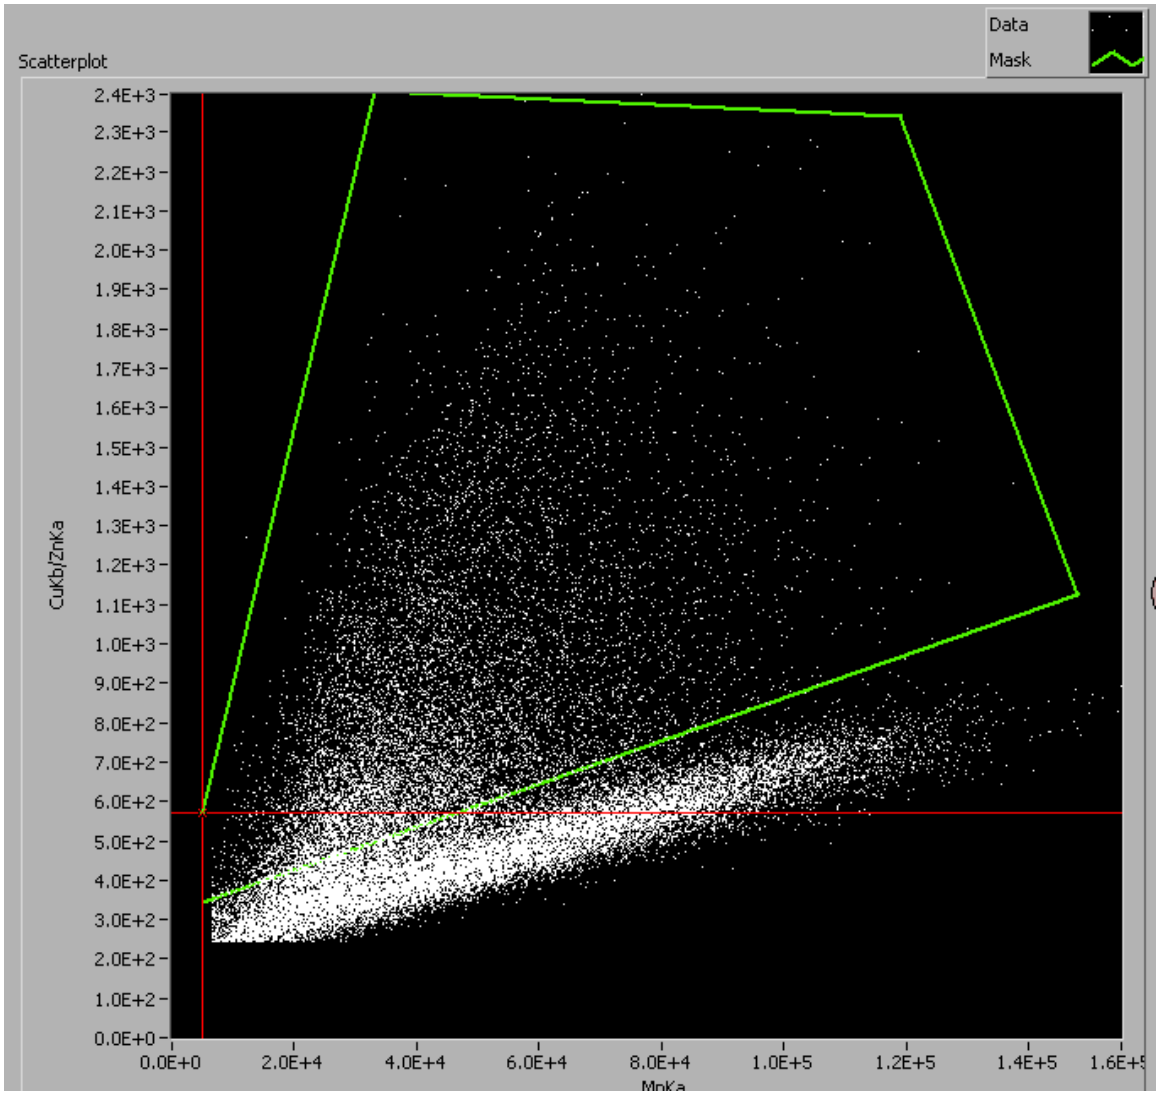

Figure 5. The previous scatterplot, but with a mask drawn to capture the upper branch. On hitting the Complete Mask button, the missing side of the polygon will be included but not drawn in on this plot.

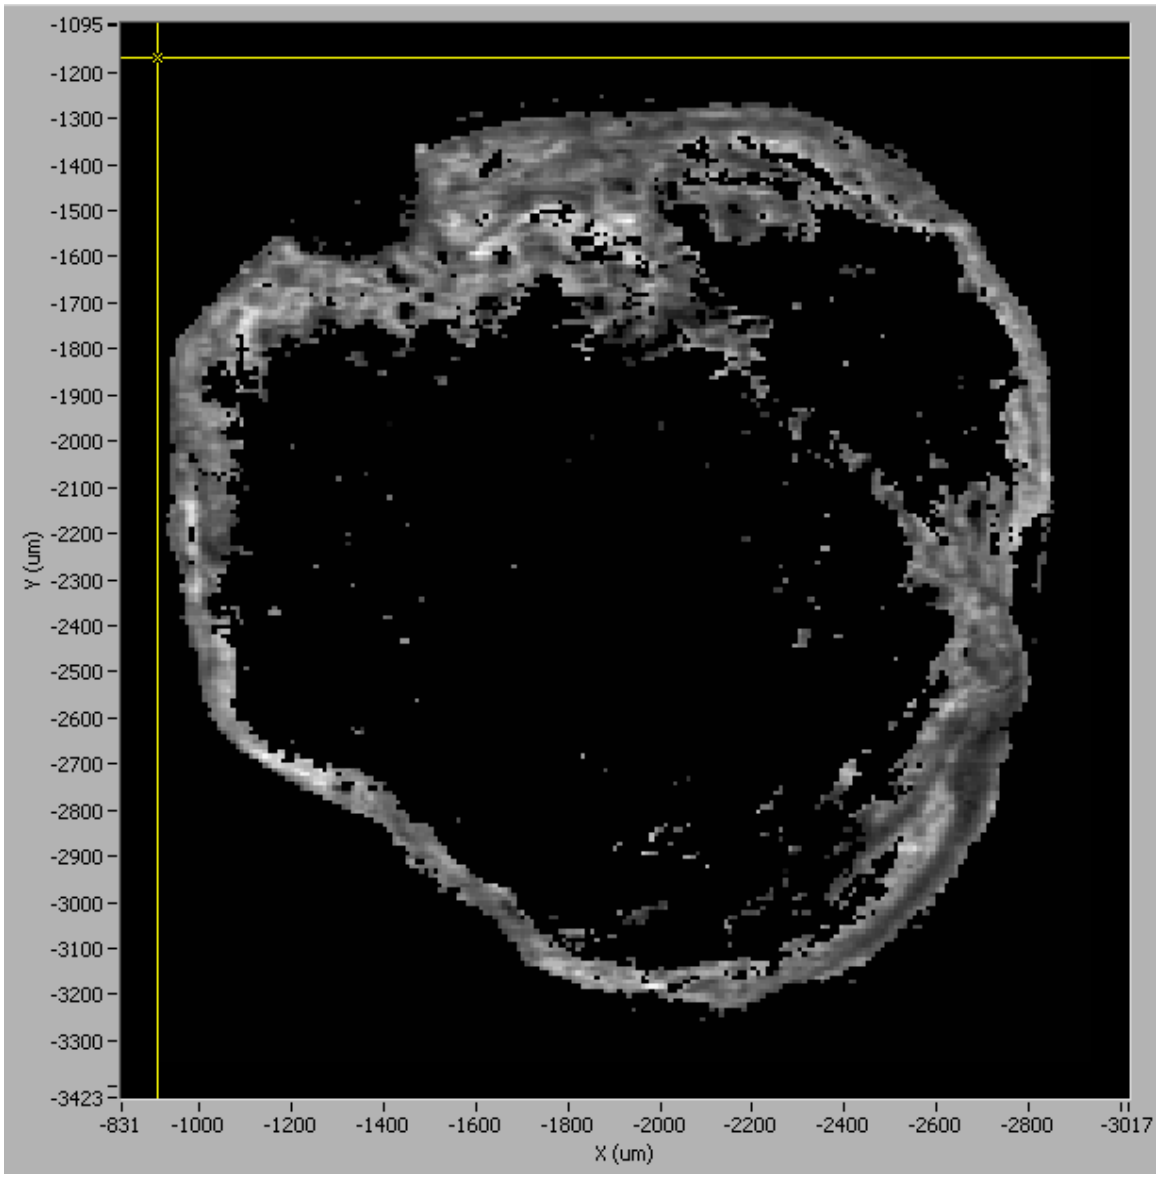

Figure 6. The Zn image, with the mask from Figure 5 applied. The mask clearly resolves a specific region in the sample.